EAST Search History

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
S1	3	lawrence near mcvoy.in.	US-PGPUB; USPAT	OR	ON	2006/10/30 09:16
S2	138	wayne near scott.in.	US-PGPUB; USPAT	OR	ON	2006/10/30 09:17
S3	1	bryan near oakley.in.	US-PGPUB; USPAT	OR	ON	2006/10/30 09:18
S4	543	717/110-113.ccls.	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/10/30 09:19
S5	435	S4 and (merg$4 or combin$5 or difference)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/10/30 09:20
S6	292	S5 and (user near interface)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/10/30 09:20
S7	16	("5278979" "5715454" "5828885" "5835601" "6226652" "6321378" "6366933" "6367077").pn.	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/10/30 10:42
S8	101	("20040177343" "20050169073" "20050172279" "5649099" "7117495" "20030233489" "20030233490" "20030233647" "5878218" "6049874" "5666501" "5765156" "5794254" "5907835" "5924099" "6014676" "4912637" "5208853" "5360651" "5495613" "5572637" "5606651" "5634114" "5774728" "5806078" "5835775" "5835773" "5845293" "5850522" "5884308" "5897640" "5909689" "5930513" "5956402" "6076111" "6086623" "6101505" "6195796" "6233589" "6269456" "6269474" "6317754" "6374250" "6397378" "6401220" "6473707" "7020573" "20050154553" "20030220985" "5288144").pn.	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/10/30 10:58

EAST Search History

S9	10	("5481722" | "5555409" | "5598208" | "5671428" | "5737740" | "5765176" | "5787280" | "5799315" | "5806078" | "5845293").PN.	US-PGPUB; USPAT; USOCR	OR	ON	2006/10/30 13:13
S10	18	("3711863" | "4595980" | "4641264" | "4706212" | "4875159" | "4912637" | "5084815" | "5181163" | "5185888" | "5191646" | "5247683" | "5280612" | "5323311" | "5325531" | "5339483" | "5504879" | "5604901" | "5671428").PN.	US-PGPUB; USPAT; USOCR	OR	ON	2006/10/30 13:48
S11	912	717/122,169,170,175.ccls.	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/10/30 14:46
S12	1608	707/203.ccls.	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/10/30 14:47
S14	387	S11 and (merg$4 or combin$4) and (display$$ or show$4)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/10/30 14:56
S16	754	S12 and (merg$4 or combin$4) and (display$$ or show$4)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/10/30 14:56
S17	430	S16 and (conflict or difference)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/10/30 14:59
S18	4	stacked adj diff	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/10/30 15:29

EAST Search History

S19	367635	(merg$4 or reconcil$4 or synchroniz$4 or combin$4) same (file or change)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/10/30 15:33
S20	97674	(merg$4 or reconcil$4 or synchroniz$4 or combin$4) with (file or change)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/10/30 15:33
S21	11278	S20 and (display$4 or show$4) and (user adj interface)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/10/30 15:34
S22	762	S21 and align$4 and blank	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/10/30 15:35
S25	97	S22 and (simultaneous$3 near (display$4 or show$4))	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/10/30 15:39
S26	89	S25 and (@pd<"20021104" or @ad<"20021104" or @prad<"20021104" or @rlad<"20021104")	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/10/30 15:40
S27	8	("5495539" | "5668897" | "5754190" | "5874951" | "5973692" | "5999182" | "6124864" | "6232983").PN.	US-PGPUB; USPAT; USOCR	OR	ON	2006/10/31 08:55
S32	584	scroll$4 near (together or "same time" or concurrent$3 or simultaneous$3)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/10/31 09:13

EAST Search History

S33	115	S32 and user adj interface	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/10/31 09:14
S34	23	(US-20030051230-$ or US-20030187705-$).did. or (US-5845293-$ or US-5806078-$ or US-5481722-$ or US-7089531-$ or US-7086046-$ or US-7017144-$ or US-6912707-$ or US-6889359-$ or US-6407753-$ or US-6233589-$ or US-6226652-$ or US-6389433-$ or US-5897640-$ or US-5893119-$ or US-5828885-$ or US-5745906-$ or US-5729743-$ or US-6738084-$ or US-6275223-$).did. or (JP-2000039947-$).did. or (GB-2355828-$).did.	US-PGPUB; USPAT; JPO; DERWENT	OR	ON	2006/10/31 09:22
S35	1	S34 and undo	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/10/31 09:23

Terms used
files merging combining reconciling synchronizing

Found **31,889** of **186,958**

Sort results by: relevance
Display results: expanded form
Save results to a Binder
Search Tips
Open results in a new window

Try an Advanced Search
Try this search in The ACM Guide

Results 1 - 20 of 200 Result page: **1** 2 3 4 5 6 7 8 9 10 next
Best 200 shown Relevance scale

1 An algebraic approach to file synchronization

Norman Ramsey, El"od Csirmaz
September 2001 **ACM SIGSOFT Software Engineering Notes , Proceedings of the 8th European software engineering conference held jointly with 9th ACM SIGSOFT international symposium on Foundations of software engineering ESEC/FSE-9**, Volume 26 Issue 5
Publisher: ACM Press
Full text available: pdf(301.78 KB) Additional Information: full citation, abstract, references, citings, index terms

A *file synchronizer* restores consistency after multiple replicas of a filesystem have been changed independently. We present an algebra for reasoning about operations on filesystems and show that it is sound and complete with respect to a simple model. The algebra enables us to specify a file-synchronization algorithm that can be combined with several different conflict-resolution policies. By contrast, previous work builds the conflict-resolution policy into the specification, or worse, ...

2 Systems and applications: A hybrid approach to optimistic file system directory tree synchronization

Tancred Lindholm, Jaakko Kangasharju, Sasu Tarkoma
June 2005 **Proceedings of the 4th ACM international workshop on Data engineering for wireless and mobile access**
Publisher: ACM Press
Full text available: pdf(220.29 KB) Additional Information: full citation, abstract, references, index terms

There are two main approaches to optimistic file system synchronization: distributed file systems and file synchronizers. The former type is characterized by a log-based approach that depends on access to file system internals, the latter by a state-based approach that utilizes the standard file system interface, which limits the efficiency of change detection.We propose a hybrid approach that 1) defines a minor extension to the semantics of the file system interface that enables efficient state ...

Keywords: XML, directory tree, optimistic synchronization, reconciliation, stackable file system, state-based

3 Models: Using the transformational approach to build a safe and generic data synchronizer

Pascal Molli, Gérald Oster, Hala Skaf-Molli, Abdessamad Imine
November 2003 **Proceedings of the 2003 international ACM SIGGROUP conference on Supporting group work**
Publisher: ACM Press
Full text available: pdf(184.01 KB) Additional Information: full citation, abstract, references, citings, index terms

Reconciliating divergent data is an important issue in concurrent engineering, mobile computing and software configuration management. Currently, a lot of synchronizers or merge tools perform reconciliations. However, they do not define what is the correctness of their synchronisation. In this paper, we propose to use a transformational approach as the basic model for reasonning about synchronisation. We propose an algorithm and specific transformation functions that realize a file system synchr ...

Keywords: operational transformation, synchronization

4 Fast detection of communication patterns in distributed executions

Thomas Kunz, Michiel F. H. Seuren

November 1997 **Proceedings of the 1997 conference of the Centre for Advanced Studies on Collaborative research**

Publisher: IBM Press

Full text available: pdf(4.21 MB) Additional Information: full citation, abstract, references, index terms

Understanding distributed applications is a tedious and difficult task. Visualizations based on process-time diagrams are often used to obtain a better understanding of the execution of the application. The visualization tool we use is Poet, an event tracer developed at the University of Waterloo. However, these diagrams are often very complex and do not provide the user with the desired overview of the application. In our experience, such tools display repeated occurrences of non-trivial commun ...

5 What is a file synchronizer?

S. Balasubramaniam, Benjamin C. Pierce

October 1998 **Proceedings of the 4th annual ACM/IEEE international conference on Mobile computing and networking**

Publisher: ACM Press

Full text available: pdf(1.21 MB) Additional Information: full citation, references, citings, index terms

6 Optimistic replication

Yasushi Saito, Marc Shapiro

March 2005 **ACM Computing Surveys (CSUR)**, Volume 37 Issue 1

Publisher: ACM Press

Full text available: pdf(656.72 KB) Additional Information: full citation, abstract, references, index terms

Data replication is a key technology in distributed systems that enables higher availability and performance. This article surveys optimistic replication algorithms. They allow replica contents to diverge in the short term to support concurrent work practices and tolerate failures in low-quality communication links. The importance of such techniques is increasing as collaboration through wide-area and mobile networks becomes popular.Optimistic replication deploys algorithms not seen in tradition ...

Keywords: Replication, disconnected operation, distributed systems, large scale systems, optimistic techniques

7 Concurrency control in advanced database applications

Naser S. Barghouti, Gail E. Kaiser

September 1991 **ACM Computing Surveys (CSUR)**, Volume 23 Issue 3

Publisher: ACM Press

Full text available: pdf(4.69 MB) Additional Information: full citation, references, citings, index terms

Keywords: advanced database applications, concurrency control, cooperative transactions, design environments, extended transaction models, long transactions, object-oriented databases, relaxing serializability

Document querying and transformation: A three-way merge for XML documents
Tancred Lindholm
October 2004 **Proceedings of the 2004 ACM symposium on Document engineering**
Publisher: ACM Press
Full text available: pdf(500.99 KB) Additional Information: full citation, abstract, references, index terms

Three-way merging is a technique that may be employed for reintegrating changes to a document in cases where multiple independently modified copies have been made. While tools for three-way merge of ASCII text files exist in the form of the ubiquitous diff and patch tools these are of limited applicability to XML documents.

We present a method for three-way merging of XML which is targeted at merging XML formats that model human-authored documents as ordered trees (e.g. rich text forma ...

Keywords: XML, collaborative editing, conflict, structured text, three-way merge

9 The LOCUS distributed operating system
Bruce Walker, Gerald Popek, Robert English, Charles Kline, Greg Thiel
October 1983 **ACM SIGOPS Operating Systems Review , Proceedings of the ninth ACM symposium on Operating systems principles SOSP '83**, Volume 17 Issue 5
Publisher: ACM Press
Full text available: pdf(1.89 MB) Additional Information: full citation, abstract, references, citings, index terms

LOCUS is a distributed operating system which supports transparent access to data through a network wide filesystem, permits automatic replication of storage, supports transparent distributed process execution, supplies a number of high reliability functions such as nested transactions, and is upward compatible with Unix. Partitioned operation of subnet's and their dynamic merge is also supported. The system has been operational for about two years at UCLA a ...

10 Industrial sessions: beyond relational tables: Coordinating backup/recovery and data consistency between database and file systems
Suparna Bhattacharya, C. Mohan, Karen W. Brannon, Inderpal Narang, Hui-I Hsiao, Mahadevan Subramanian
June 2002 **Proceedings of the 2002 ACM SIGMOD international conference on Management of data SIGMOD '02**
Publisher: ACM Press
Full text available: pdf(1.44 MB) Additional Information: full citation, abstract, references, index terms

Managing a combined store consisting of database data and file data in a robust and consistent manner is a challenge for database systems and content management systems. In such a hybrid system, images, videos, engineering drawings, etc. are stored as files on a file server while meta-data referencing/indexing such files is created and stored in a relational database to take advantage of efficient search. In this paper we describe solutions for two potentially problematic aspects of such a data ...

Keywords: DB2, content management, database backup, database recovery, datalinks

11 External memory algorithms and data structures: dealing with **massive data**
Jeffrey Scott Vitter
June 2001 **ACM Computing Surveys (CSUR)**, Volume 33 Issue 2
Publisher: ACM Press
Full text available: pdf(828.46 KB) Additional Information: full citation, abstract, references, citings, index terms

Data sets in large applications are often too massive to fit completely inside the computers internal memory. The resulting input/output communication (or I/O) between fast internal memory and slower external memory (such as disks) can be a major performance bottleneck. In this article we survey the state of the art in the design and analysis of external memory (or EM) algorithms and data structures, where the goal is to exploit locality in order to reduce the I/O costs. We consider a varie ...

Keywords: B-tree, I/O, batched, block, disk, dynamic, extendible hashing, external memory, hierarchical memory, multidimensional access methods, multilevel memory, online, out-of-core, secondary storage, sorting

12 Reconciling environment integration and software evolution

Kevin J. Sullivan, David Notkin
July 1992 **ACM Transactions on Software Engineering and Methodology (TOSEM)**, Volume 1 Issue 3
Publisher: ACM Press

Full text available: pdf(2.89 MB) Additional Information: full citation, abstract, references, citings, index terms, review

Common software design approaches complicate both tool integration and software evolution when applied in the development of integrated environments. We illustrate this by tracing the evolution of three different designs for a simple integrated environment as representative changes are made to the requirements. We present an approach that eases integration and evolution by preserving tool independence in the face of integration. We design tool integration relationships as separate component ...

Keywords: abstract behavior type, behavior abstraction, component independence, environment integration, event mechanism, implicit invocation, integrated environment, mediator, mediator/event design, software evolution, tool integration

13 Impact of software engineering research on the practice of software configuration management

Jacky Estublier, David Leblang, André van der Hoek, Reidar Conradi, Geoffrey Clemm, Walter Tichy, Darcy Wiborg-Weber
October 2005 **ACM Transactions on Software Engineering and Methodology (TOSEM)**, Volume 14 Issue 4
Publisher: ACM Press
Full text available: pdf(350.59 KB) Additional Information: full citation, abstract, references, index terms

Software Configuration Management (SCM) is an important discipline in professional software development and maintenance. The importance of SCM has increased as programs have become larger, more long lasting, and more mission and life critical. This article discusses the evolution of SCM technology from the early days of software development to the present, with a particular emphasis on the impact that university and industrial research has had along the way. Based on an analysis of the publicati ...

Keywords: Versioning, data model, process support, research impact, software configuration management, software engineering, workspace management

14 Process model and awareness in SCM

Jacky Estublier, Sergio Garcia
September 2005 **Proceedings of the 12th international workshop on Software configuration management SCM '05**
Publisher: ACM Press
Full text available: pdf(565.84 KB) Additional Information: full citation, abstract, references

The development of large and complex systems, under hard time constraints, requires the participation of many developers working concurrently. SCM systems allow concurrent access to software artifacts, but provide poor support to maintain data consistency when concurrent changes are performed on the same artifacts. This problem can be reduced if developers are aware of the others work and warned about the conflicts that may arise, allowing the users to manage the risks more effectively.Awareness ...

15 Multi-level transaction management for complex objects: implementation, performance, parallelism

Gerhard Weikum, Christof Hasse
October 1993 **The VLDB Journal — The International Journal on Very Large Data**

Bases, Volume 2 Issue 4
Publisher: Springer-Verlag New York, Inc.
Full text available: pdf(2.83 MB) Additional Information: full citation, abstract, references, citings

Multi-level transactions are a variant of open-nested transactions in which the subtransactions correspond to operations at different levels of a layered system architecture. They allow the exploitation of semantics of high-level operations to increase concurrency. As a consequence, undoing a transaction requires compensation of completed subtransactions. In addition, multi-level recovery methods must take into consideration that high-level operations are not necessarily atomic if multiple pages ...

Keywords: atomicity, complex objects, inter- and intratransaction parallelism, multi-level transactions, performance, persistence, recovery

16 Concepts in configuration management systems

Susan Dart
May 1991 **Proceedings of the 3rd international workshop on Software configuration management**
Publisher: ACM Press
Full text available: pdf(1.92 MB) Additional Information: full citation, references, citings, index terms

17 Distributed file systems: concepts and examples

Eliezer Levy, Abraham Silberschatz
December 1990 **ACM Computing Surveys (CSUR)**, Volume 22 Issue 4
Publisher: ACM Press
Full text available: pdf(5.33 MB) Additional Information: full citation, abstract, references, citings, index terms, review

The purpose of a distributed file system (DFS) is to allow users of physically distributed computers to share data and storage resources by using a common file system. A typical configuration for a DFS is a collection of workstations and mainframes connected by a local area network (LAN). A DFS is implemented as part of the operating system of each of the connected computers. This paper establishes a viewpoint that emphasizes the dispersed structure and decentralization of both data and con ...

18 Using structural characteristics for autonomous operation

Carlos Baquero, Francisco Moura
October 1999 **ACM SIGOPS Operating Systems Review**, Volume 33 Issue 4
Publisher: ACM Press
Full text available: pdf(636.87 KB) Additional Information: full citation, abstract, index terms

The majority of current mobile computing systems operate either in conjunction with a central network by some form of weak connectivity or tend to operate in total isolation and perform sporadic synchronization with a backup or a central network. These configurations miss an additional and very useful pattern of operation --- mobile to mobile interaction. Recent mobile devices have the capacity for direct communication among them, but this option is essentially neglected by the application softw ...

Keywords: conflict resolution, mobile computing, replication

19 Client-server computing in mobile environments

Jin Jing, Abdelsalam Sumi Helal, Ahmed Elmagarmid
June 1999 **ACM Computing Surveys (CSUR)**, Volume 31 Issue 2
Publisher: ACM Press
Full text available: pdf(233.31 KB) Additional Information: full citation, abstract, references, citings, index terms, review

Recent advances in wireless data networking and portable information appliances have engendered a new paradigm of computing, called mobile computing, in which users carrying portable devices have access to data and information services regardless of their

physical location or movement behavior. In the meantime, research addressing information access in mobile environments has proliferated. In this survey, we provide a concrete framework and categorization of the various way ...

Keywords: application adaptation, cache invalidation, caching, client/server, data dissemination, disconnected operation, mobile applications, mobile client/server, mobile compuing, mobile data, mobility awareness, survey, system application

20 The IceCube approach to the reconciliation of divergent replicas

Anne-Marie Kermarrec, Antony Rowstron, Marc Shapiro, Peter Druschel
August 2001 **Proceedings of the twentieth annual ACM symposium on Principles of distributed computing**
Publisher: ACM Press

Full text available: pdf(751.74 KB) Additional Information: full citation, abstract, references, citings, index terms

We describe a novel approach to log-based reconciliation called IceCube. It is general and is parameterised by application and object semantics. IceCube considers more flexible orderings and is designed to ease the burden of reconciliation on the application programmers. IceCube captures the static and dynamic reconciliation constraints between all pairs of actions, proposes schedules that satisfy the static constraints, and validates them against the dynamic constraints.

Preliminary ...